

In the Specification

Please replace paragraph [0012] with the following new paragraph:

[0012] Referring now to Figure 1, a schematic diagram 100 of portions of a pipeline of a processor is shown, according to one embodiment. The Figure 1 pipeline may include a front end 120 and a back end 140. The front end 120 may include a fetch/decode stage 122. The fetch/decode stage 122 may fetch (and potentially prefetch) instructions from a memory hierarchy that may in one embodiment include an L1 cache 106, an L2 cache 104, and memory 102. In other embodiments, other forms of memory hierarchy may be used. The fetch/decode stage 122 may also decode the instructions. In one embodiment, macro-instructions may be decoded into sets of micro-operations.

Please replace paragraph [0017] with the following new paragraph:

[0017] Referring now to Figure 2, a schematic diagram of portions of a pipeline of a processor including two register alias tables is shown, according to one embodiment. In the Figure 2 embodiment a trace cache 210 is shown, but in other embodiments other forms of instruction buffers may be used to store pending instructions, including micro-operations. Micro-operations supplied by the trace cache 210 send their logical register addresses to the RAT 220 for translation to physical register addresses. As shown in Figure 2, re-scheduler 240 is coupled to RAT 220. In processors conforming to the Intel® Corporation Pentium® architecture, each micro-operation may require up to 2 source operands along with 1 destination operand. A read port may be required for each source operand, and an additional read port may be required to read out the old mapping for the destination register. (The old mapping may be read before it is overwritten in order to free it when the micro-operation is retired.) Consider a back end capable of executing N

instructions in parallel. Then a worst-case scenario would require that the RAT 220 have 3 read ports per instruction or a total of $3N$ read ports. Similarly, a destination operand may also require a write port to the RAT 220, so then a worst-case scenario would require that the RAT 220 have 1 write port per instruction or a total of N write ports. Using a conventional RAT, this would mean that a total of $3N$ read ports and N write ports for each and every logical register address. It is noteworthy that the physical size of a multi-ported structure, such as a RAT, typically grows in size in proportion to the square of the number of ports. For example, a structure with $4N$ ports ($3N + N$ ports) may be 16 times larger than an equivalent structure that only has N ports. Thus, as the number of ports increase, the structures may become increasingly expensive in terms of physical size, power consumption, and time to design.

Please replace paragraph [0022] with the following new paragraph:

[0022] In another technique, an architecture may include several “control” registers that are used infrequently. The “control” registers may contain information that determines certain modes of operation of the machine. Some examples of such modes may include how the rounding operations are performed after floating point operations, how the floating point exceptions are handled, and how the floating point denormal operations are handled. The information for the mode may be kept in a set of control registers that are infrequently changed. Such control registers may be ~~may be determined~~ to be candidates for being serviced by a low-bandwidth RAT 224. In yet another technique, a given compiler or set of compilers may tend to use some registers much more often than others. If so, then these compiler artifacts may be examined to determine which registers may be candidates for being serviced by a low-bandwidth RAT 224. In some embodiments, the selected registers may be used as a requirement on compiler designers.

Please replace paragraph [0025] with the following new paragraph:

[0025] Referring now to Figure 3, a schematic diagram of traces in a trace cache is shown, according to one embodiment of the present disclosure. Figure 3 illustrates traces 340, 360 of one embodiment, in the context of a 4-way (330-336), 256-set (310-320, ...324) embodiment of trace cache array 300 storing instructions that are decoded micro-ops of macro-instructions. Each storage location in the array, called a data line, may be identifiable by a set number and way number. Based on the descriptions to follow, those skilled in the art will appreciate that other sizes of set-associate caches, as well as non-set-associate caches, may be employed in other embodiments. Moreover, the instructions may be instructions of any kind.

Please replace paragraph [0026] with the following new paragraph:

[0026] For ease of further discussion, the first trace element of a trace 340 may be referred to as a trace head, or head 342; the last trace element of a trace being referred to as a trace tail, or tail 352; and the intermediate trace elements being referred to as trace bodies, or bodies 344, 346, 348, and 350. As shown in Figure 3, trace 360 has head 362, bodies 364 and 366, and tail 368. In the degenerate case of a two-element trace, the second trace element is a trace body as well as a trace tail, and in the degenerate case of a single element trace, the singular trace element is a trace head, a trace segment body, as well as a trace tail at the same time.

Please replace paragraph [0032] with the following new paragraph:

[0032] Figure 4 shows a RAT 418 that includes three component RATs: a high-bandwidth RAT 422, a mid-bandwidth RAT 424, and a low-bandwidth RAT 426. As

shown in Figure 4, trace cache 410 is coupled to RAT 418. Re-scheduler 440 is coupled to RAT 418, as shown in Figure 4. The high-bandwidth RAT 422 and low-bandwidth RAT 426 are shown with the corresponding number of read ports 428, 436, respectively, and write ports 430, 438, respectively, as used by the high-bandwidth RAT 222 and low-bandwidth RAT 224 of Figure 2. However, in other embodiments other numbers of read ports and write ports may be used. The mid-bandwidth RAT 424 may have a number of read ports and a number of write ports somewhere between that used by the high-bandwidth RAT 422 and the low-bandwidth RAT 426. In the Figure 4 embodiment, mid-bandwidth RAT 424 is shown with N read ports and N/2 write ports, although other numbers could be chosen.

Please replace paragraph [0034] with the following new paragraph:

[0034] Referring now to Figure 5, a schematic diagram of a microprocessor system is shown, according to one embodiment of the present disclosure. The Figure 5 system may include several processors of which only two, processors 40, 60 are shown for clarity. Processors 40, 60 may include level one caches 42, 62. The Figure 5 multiprocessor system may have several functions connected via bus interfaces 44, 64, 12, 8 with a system bus 6. As shown in Figure 5, the connection of processors 40 and 60 to system bus 6 is illustrated by lines 48 and 68 respectively. In one embodiment, system bus 6 may be the front side bus (FSB) utilized with Pentium® class microprocessors manufactured by Intel® Corporation. In other embodiments, other busses may be used, or point-to-point interfaces may be used instead of a bus. A general name for a function connected via a bus interface with a system bus is an "agent". Examples of agents are processors 40, 60, bus bridge 32, and memory controller 34. In some embodiments memory controller 34 and bus bridge 32 may collectively be referred to as a chipset. In

some embodiments, functions of a chipset may be divided among physical chips differently than as shown in the Figure 5 embodiment.